

A Testbed for Data Fusion for Helicopter Diagnostics and Prognostics¹

Tom Brotherton and Paul Grabill
IAC - The Intelligent Automation Corp
Poway, CA 92064
(858) 679-4140
tom.brotherton@iac-online.com

Richard Friend and Bill Sotomayer
Aeronautical Systems
Wright-Patterson AFB, OH 43558

John Berry
US Army Aviation and Missile Command
Redstone Arsenal, AL 35898
(256) 313-4815
john.berry@redstone.army.mil

Abstract— Automated systems to perform aircraft diagnostics and prognostics are of current interest. Development of those systems requires large amounts of data (collection, monitoring, manipulation) to capture and characterize normal, known fault events, and to ensure data is captured early on in a fault progression to support prognostic system development. Continuous data collection is also required to capture relatively rare events. Data collected can then be analyzed to assist in the development of automated systems and for continuous updating of algorithms to improve detection, classification, and prognostic performance. IAC, in collaboration with the Air Force and Army, is developing a testbed on which to perform data collection, and develop diagnostic and prognostic processing techniques using Army helicopter vibration and engine performance data as part of the Army's Vibration Management Enhancement Program (VMEP). VMEP and the testbed being developed for collection and processing of VMEP data are described here.

TABLE OF CONTENTS

1. INTRODUCTION
2. OVERVIEW OF VMEP
3. VMEP WEB SERVER
4. VMEP SERVER OPERATION - EXAMPLES
5. SUMMARY AND RECOMMENDATIONS

1. INTRODUCTION

The Intelligent Automation Corporation (IAC), in collaboration with the US Army and Air Force is developing a system for monitoring of Army helicopter vibration, engine performance, and aircraft structures. The processing being developed is part of the Army's Vibration Management Enhancement Program (VMEP).

VMEP is composed of three primary components. The first component is an on-board system that measures and processes vibration and parameter information in flight. The second component is a ground-based software system that displays recommended maintenance actions at the aircraft, aircraft status to the maintenance manager, and

measurement details to the engineer. The third component is a system of web-based tools that provides data archiving, software configuration management, management reports, and an advanced engineering development testbed.

IAC supports development all three components of VMEP. In this paper we will focus on the third component; the VMEP Server. This system includes an Internet utility to collect data from the ground-based software located at the unit/aircraft for higher-level comparisons and statistical modeling as well as update of the algorithms and parameters of aircraft on-board systems. All of the tools have been developed using standard web server development software as well as the Mathwork's Matlab and Simulink tools.

Recently, there has been a lot of hype in "prognosis." Diagnostics and prognostic problems are similar in nature [1]. Both are looking for patterns that are indications of faults. Once a fault pattern is recognized, then something needs to be done about it. In diagnostic problems, the "signal-to-noise ratio" (SNR) of the fault signals to the ambient background is high; the fault is well developed and its signature is easily seen. And by definition, a fault has already occurred, so that some immediate maintenance action needs to be taken. For prognosis, the equivalent fault SNR is much lower so that the fault signature is hard to pick out of the background. Since by definition of prognosis, nothing is immediately wrong, the problem becomes prediction of the time horizon before something needs to be done.

Because faults in the prognosis problem appear at low SNR, they are hard to see. Making accurate predictions, while maintaining an acceptable false alarm rate, is much harder to do then with the diagnostics problem. One solution is to integrate (or fuse) low level signals and information that may be seen across a variety of sensors or algorithms to effectively increase the fault SNR.

There is a problem in "advanced" prognosis where there does not exist good data sets to support its development.

¹ Paper # 1364 - This effort is being funded by USAF and US Army
0-xxxx-xxxx-X/01/\$10.00/©2003 IEEE
U.S. Government work not protected by U.S. copyright.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 01 MAR 2003		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE A Testbed for Data Fusion for Helicopter Diagnostics and Prognostics1				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) IAC - The Intelligent Automation Corp Poway, CA 92064				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES Proceedings of the 2003 IEEE Aerospace Conference, Big Sky MT, March 2003					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 14	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

All the action in prognosis occurs in the tails of distributions of “normal” vs. “fault” measurements. However it is in these tails that no data has ever or is just now starting to be collected. This data is required for empirical model development but also for validation of detailed physics based models. In current data sets there is lots of “normal” and “fault” (easily detected) data near the mean of “normal” and “fault”. The tails are a scary place. A slight change in a threshold can mean a drastic change in false alarms, missed detections, and prognostic prediction time horizons.

The VMEP system and in particular the web component are ideal for performing data collection and algorithm design and tuning in order to develop advanced diagnostic and prognostic techniques for air craft health monitoring.

Here a description of the overall VMEP system will be given with emphasis on the Web Server component. A description of some of the tools that have been developed and the results of their application to processing real Army helicopter data are presented.

2. OVERVIEW OF THE VIBRATION MANAGEMENT ENHANCEMENT PROGRAM

The Vibration Management Enhancement Program (VMEP) is a helicopter vibration and health monitoring system.

The primary function of the VMEP system is to provide a built-in capability to perform routine vibration maintenance functions (such as rotor smoothing and mandatory vibration checks) during routine operational flights. In addition, the system monitors the status or health of the dynamic drive system components and engine related exceedances. A capability for flight regime recognition / structures monitoring is currently being added. The availability of advanced signal processing for machinery fault diagnostics allows much of the processing of vibration signatures and other monitoring operations to be completed during in-flight operation of the aircraft. The VMEP is intended to detect faults with sufficient lead time so that the ground-maintainer can schedule corrective actions well before the fault becomes an in-flight failure.

Overall system description

The VMEP system is composed of three primary components. The first component is an on-board system which measures and processes vibration and parameter information in flight. The second component is a ground-based software system which displays recommended maintenance actions at the aircraft, aircraft status to the maintenance manager, or measurement details to the engineer. The third component is a web server which




Figure 1 The VMEP on-board system

provides data archiving, software configuration management, aircraft status reports, and an advanced engineering development testbed. Details of the overall VMEP system can be found in [2,3].

On-board system

The on-board system, shown in Figure 1, consists of a Vibration Management Unit (VMU), a wiring harness, and sensors. The VMU front panel provides the aircrew a method of selecting acquisitions at specific flight conditions and receiving system status information. The sensors include tachometers and accelerometers distributed throughout the helicopter's drive train.

The data acquisition process on board the VMU is configurable based on the type of flight. The system can be setup for engineering data acquisition or for day-to-day data collection. An engineering setup may include collecting data in a raw format like a digital tape recorder. This allows for the most flexibility in post processing. In normal day-to-day operation the VMEP is setup to pre-process the data and only store condensed Condition Indicators (CIs) in small compact data files.

If a new problem is found in a mechanical component, a small change in the setup file can be made to allow the VMU to collect raw or intermediate results for detailed engineering analysis.

The data that is collected and processed in the VMU is stored for data transfer after the aircraft lands. The current VMU has 96 Mbytes of non-volatile memory for program and data storage. A typical flight contains fewer than 100 Kbytes of data allowing 900 typical flights to be stored before the data needs to be downloaded. The typical operation at a facility has the operators downloading the flights data at the end of the day. The typical download process only takes one minute. The size of the data files can be changed if engineering desires more raw data and less pre-processed Condition Indicators.

Ground based station

The ground-based software runs on a PC based Windows platform. The system is referred to as the PC Ground-Based System (PC-GBS). Figure 2 shows the opening screen and sample selection of summary reports available from the PC-GBS. The operator downloads the processed data from the VMU after data has been captured in flight

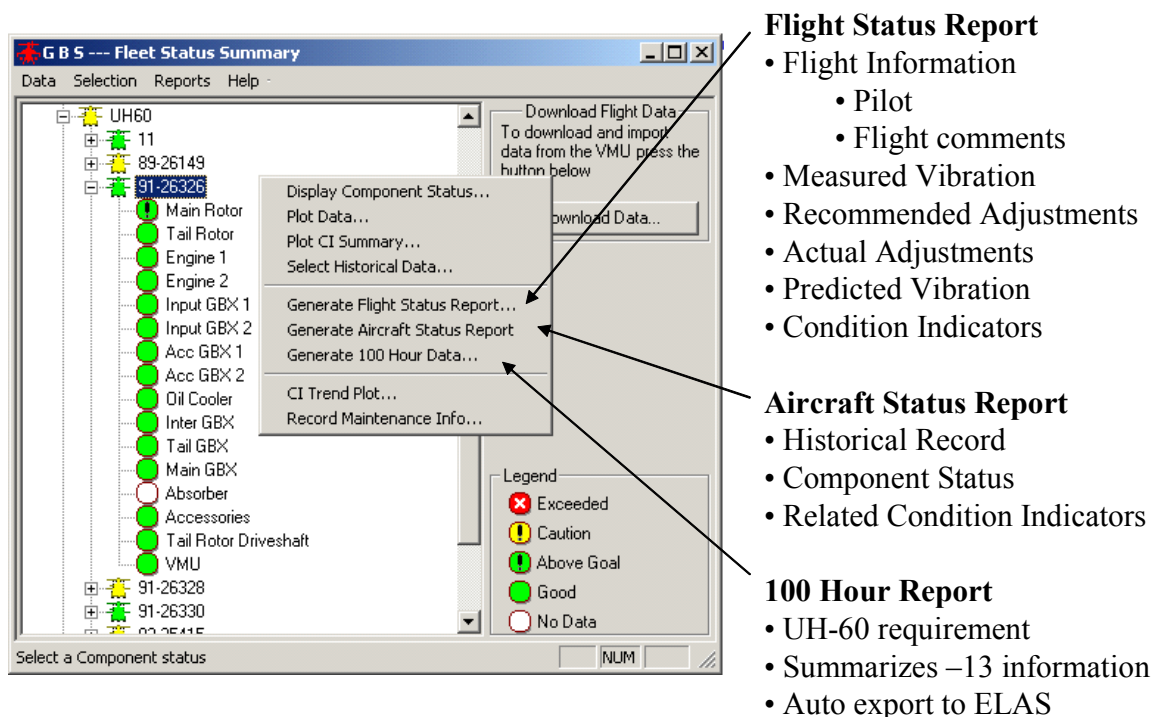


Figure 2 VMEP PC-Ground Based Station – Example of summary report selection

via a serial cable. This software interprets the processed data and provides a multi-level operator interface that is oriented to provide specific data to assist skilled maintainers in isolating potential faults.

Where sufficient data is known about a specific fault indicator, the instructions are provided for corrective action. This software also allows the aggregation of multiple aircraft for comparisons and fleet statistics by maintenance managers. The software also allows detailed examination of the data by engineering personnel.

More importantly for the work presented here, the on-board system and PC-GBS are a continuous data collection system. This data is automatically downloaded to the web component of VMEP when ever the PC-GBS is attached to the web system.

3. VMEP WEB SERVER

Connectivity

The VMEP system includes an internet utility to collect, analyze, and make available data from the PC-GBS located at the unit/aircraft. The successful development of algorithms for helicopter condition health monitoring requires real data that represent specific conditions.

These conditions are: nominal operation; operation with known faults; and, most importantly for prognostics, operation leading up to the time that a fault can be detected. Typically, data is saved only when a fault is detected; too late to be useful for prognostics development.

Figure 3 shows the hardware and connectivity of the system used to perform data collection. The VMU collects vibration and engine performance data. Maintenance personnel download the data, via serial port, to a PC-GBS which is typically a ruggedized laptop computer

Data is then transferred to the VMEP Server automatically when the PC-GBS and server are connected. Agent based software detects if new files exist on the PC-GBS that do not exist on the server. If not, that data is automatically sent.

Users that do not have the PC-GBS can also have access to data and results using a standard web browser interface.

Functional flow of VMEP Server

Figure 4 shows an overview of the functional flow of the




Figure 3 VMEP system connectivity




Figure 4 VMEP Server functional overview

VMEP Server system. There are two major components to that system as indicated in the figure. The first is the *VMEP Server* itself. As described above, the VMEP Server is the central data repository for collection and organization of all aircraft collected data. This processing is handled by the Automated Data Archive in the Server.

Network Security is performed using a secure internet connection and password protection to access the system. The user / password are context sensitive so that user's will only be able to access those components of the web system for which they are authorized.

Aircraft Configuration and Software Updates for the PC-GBS and VMU are stored on the VMEP Server and automatically transferred during WEB connection using Configuration Control. The updates are first transferred to the PC-GBS and then onto individual aircraft on-board systems as opportunities arise (i.e. when maintenance personnel interface a PC-GBS to a specific tail number VMU). Configuration control will inform maintenance if too much time has passed between the time an upgrade is posted and it has not been transferred to a specific tail number. All the maintenance needs to do is attach a PC-GBS with the updates to the on-board system that does

not have the updates.

A most useful portion of the Server for maintenance is the Fleet Statistics & Reports section. Here, fleet data, statistics, trending, and summary reports are available. These reports are available down to individual tail number and component level.

Electronic 'Help' is available in the form of electronic user's manuals, power point training presentations, and FAQs.

Advanced engineering contains a variety of modules to analyze the incoming data. These include Diagnostics, Prognostics, and Novelty Detection.

Diagnostics and prognostics algorithms are designed to respond to known fault conditions. A novel event is an unknown off-nominal condition. That is, the novel event is not nominal nor is it classified in any of the known fault conditions. *It's something completely new.*

Novelty detection is an important component in the operation of the Server. All incoming data is screened to detect, set aside, and flag for engineering analysis

anomaly events. Engineers will not have to continuously examine “normal” events. Rather only “interesting” events need be examined. It is these sorts of events that are on the edges of the data distributions between normal and fault that are of the most interest for developing prognostic algorithms.

The second component of Server is the *intelligent Machinery Diagnostics System* (iMDS) Development system. The iMDS Development system is a “behind the scenes” set of tools used by diagnostic engineers. It contains tools for performing advanced engineering analysis on data stored on the Server and elsewhere. The toolkit allows engineers to prototype algorithms that can later be incorporated into upgrades for the PC-GBS and on-board systems.

iMDS Development is standalone from the other VMEP Server system components; however, it has the ability to download and process data from the iMDS Server. Details of the iMDS Development system can be found in [1, 2].

4. VMEP SERVER OPERATION - EXAMPLES

Figure 5 shows the opening screen the user sees when entering the VMEP Server via the browser interface. There are 5 major links from the home page.

Fleet Analysis is designed for the maintainer and maintainer support personnel. It contains graphical summaries of fleet status as well as details of individual aircraft, aircraft component, and individual condition indicator (CI) status. Configuration control of VMEP software releases is also included.

Advanced Engineering is designed for engineers. It allows for visualization of data sets, selection and labeling of ‘normal’ and ‘fault’ representative data sets, setting of individual component CI detection thresholds, and development of models for diagnostics, prognostics, and anomaly detection.




Figure 5 VMEP Server Browser Interface

Download contains the latest releases of both the PC-GBS and on-board system software as well as archived data and the latest VMEP related user manuals.

Support contains links to FAQ's, a bug reporting mechanism, and aircraft (non-VMEP) related user's manuals.

Related Links contain links to other websites that may be of interest to the user.

Fleet Analysis

Whenever possible, visualization of data processing results and summaries have been used in the Server. The browser interface uses all the standard pull down menus, 'back' button, and hyperlinks familiar to users. *'Fleet Analysis'* is the main summary page used by maintenance and fleet maintenance support personnel. Figure 6 shows a sample screen that is available from *Fleet Overview*.




Figure 6 Fleet Overview example

The figure shows at the top a summary status for all AH64 aircraft at a unit; of the 21 aircraft being monitored 24% have at least one component that has a red (in exceedance) status, 52% are Yellow (or caution) status, and 24% are Green (good) status. The numbers / percentages shown here are for demonstration purposes only. Note this is not a complete data set, rather selected tails / times to give a good demonstration.

Details on which are the troublesome components can be

found in the bottom portion of the display. Here we see that the Tail Rotor is the most troublesome with 15% of the aircraft having tail rotor vibrations in exceedance of the specified levels. Additional details can be found by clicking on the Tail Rotor hypertext to bring up all of the CIs computed for the tail rotor. A similar bar graph summary display is presented. The CI level information further isolates problematic areas.

The Fleet Analysis window also allow maintenance and maintenance support to quickly see what the readiness of all / or specific type aircraft at a unit or across the fleet is. The most troublesome components and the faults associated with them can quickly be identified.

Figure 7 shows an example of the detail available when *Aircraft Status by Tail Number* is selected. A tree structure similar to a Windows Explorer tree is brought up. That tree can be expanded / collapsed to supply the user with the detailed required.




Figure 7 Aircraft Status example

Red, yellow, and green colors are again used in the icons to indicate the current status of all the aircraft at a given point in the tree. Green means everything is within tolerance. Red means that some component is out of tolerance. Yellow indicates that the component is in a 'warning' band and requires close monitoring. For example, in Figure 7 at least one of the AH64s has a red status while at least one of the UH60 has yellow status. The tree has been expanded to show that for a particular UH60 tail number, the component that leads to the caution status is due to the Main Rotor.

Advanced Engineering

Advanced Engineering is intended for the more sophisticated engineering user. Currently, *Advanced Engineering* includes:

- Database statistics
- Fleet statistics and
- Anomaly detection.

Figure 8 shows an example of the output obtained when the user selects *Database Statistics* from the *Advanced Engineering* window. The example shows all the data that has been stored for all AH64 aircraft to date. As shown, a total of 857 data sets has been collected which contain 8750 CIs that have been found from that data. This is really not nearly enough if we want to use real data to specify a 10^{-5} false alarm threshold.

Aircraft type: AH64

Mode	Number of Data Sets	Number of CIs
BIT	2532	2532
FLIGHT	518	23545
GNDBAL	157	142
MONITOR	1583	11023
SURVEY	136	47
TAIL	166	155
TRACK	178	168
Totals	5270	37612

Figure 8 Database statistics for AH-64




Figure 9 Advanced Engineering Fleet Statistics example

Fleet Statistics brings up a display that compares single component – single CI for a particular tail number to all aircraft within the fleet. Similar to the Aircraft Status summary tree, Fleet Statistics summary for all the CIs can be obtained through a similar tree.

On Figure 9, the left side shows the Fleet Statistics summary / selection tree expanded to highlight specific tail number 86-08996. For that tail, the current condition is Red (there is a component in exceedance that is not shown). The individual components indicate that Engine 2 vibrations are in the Yellow / caution limit and that the CI labeled *SP2 FPG100 #2 Eng GG* gives rise to the caution.

The right hand side of Figure 9 shows a *box plot* summary of the selected tail / component / condition indicator compared to that CI found for all AH64 aircraft.

A box plot summarizes the CI data. It indicates the median, upper and lower *quartile*, upper and lower *adjacent values*, and *outlier* individual points. The boxplot in Figure 9 is broken down as follows. The large dot in the plot shows the median (middle point value) of the data. The ‘box’ is drawn so that 50% of the data will reside inside the box. 25% of the data is to the left of the box and 25% is to the right of the box. The dotted lines are called *fences* or *gates*. The gates are sort of a poor man’s single CI anomaly detector; if the data is well behaved all the points should fall between the gates. Outlier points are plotted as individual circles outside the gates.

Aircraft Type: AH64
Tail Number: 86-08996
Component: Engine 2




Figure 10 Anomaly Detection example

In Figure 9 we see that the points associated with the selected tail number are self consistent. For the most part (with the exception of a 0 value) they all lie within the blue box plus gates. However it is easily seen in the display that as a group, they are very different from the same CI calculated from all the other aircraft in the database.

Anomaly Detection

Up to this point, each of the CIs has been handled as single variables. Fusing information from several CIs will take advantage of the relationships of CIs in order to improve component level processing.

There are a variety of anomaly detectors (ADs) that can be used for the problem [4]. On the Server we use a neural net solution to the problem [5]. The basic neural

net anomaly detector uses radial basis function (RBF) neural nets to form a statistical model of “nominal” data. As new data enters into the system, it is compared to the RBF neural net model. If data falls within the boundaries defined by that model, then it is flagged as “nominal”. If it does not, then it is flagged as an “anomaly”. In the web site several different anomaly detectors are implemented using not only the RBF neural network, but *Support Vector Machines* and *Fuzzy Logic*. The user can call up the different detectors from a pull down window.

Notice that there is an implicit Gaussian assumption in the AD as implemented using the RBF neural network. In order to visualize distribution of the multi-CI data we have created sets of 2-dimensional plots (one for each pair of CIs) to see how the CIs are related and to determine if the model fit is appropriate. Figure 10 shows an example of that plot for the Engine #2 on the AH64

aircraft.

The Engine #2 component has 3 CIs associated with it. Thus there is the 3 x 3 array of plots shown in Figure 10. The plots correspond to pairs of CI values for the off diagonal plots. The plots along the main diagonal are histograms for the corresponding CI. Note that the scales on the plots are different depending on what is being plotted. The CI values plotted are labeled at the top and sides.

The CI that was flagged in the single CI case corresponded to tail number 86-08996. Those points stand out as a cluster in two dimensions in the figure. For anomaly detection purposes, those points were labeled as “good” along with all the other green points shown in the plot. A neural net model was built using those green points. Two outliers, identified as the red points, were found by the anomaly detector.

Single CI Detection Threshold Setting

Initial settings for single CIs on the Server have either been set based on the original manufacturers specifications or by engineering judgment as to what is acceptable or not. Since a substantial amount of data for Army aircraft has been collected, the website has included an automated statistics based setting for the various thresholds of single CIs. The automated setting of thresholds is based on ordered statistics of the data.

Figure 11 shows an example of that page for the AH64 Drive shafts for single CI threshold setting. There are four CIs associated with that component but only two are shown in the figure (the user would scroll down on the website for the other two). The current and suggested values of the ‘goal’ (green), ‘caution’ (yellow), and ‘exceedance’ (red) are shown.

The *Recommended Settings* are found directly from the collected normal data using the box plot processing similar to that described before. Figure 11 indicates that the Recommended Settings for the first CI (SP2 FPG100 Aft HB Drive Shaft), are considerably above the current settings. Indeed this particular CI was a problem because of the low threshold settings which created numerous false alarms.

As seen in Figure 11, the user has the option to accept each of the recommended settings, fill in (and accept) their own settings, or do nothing.

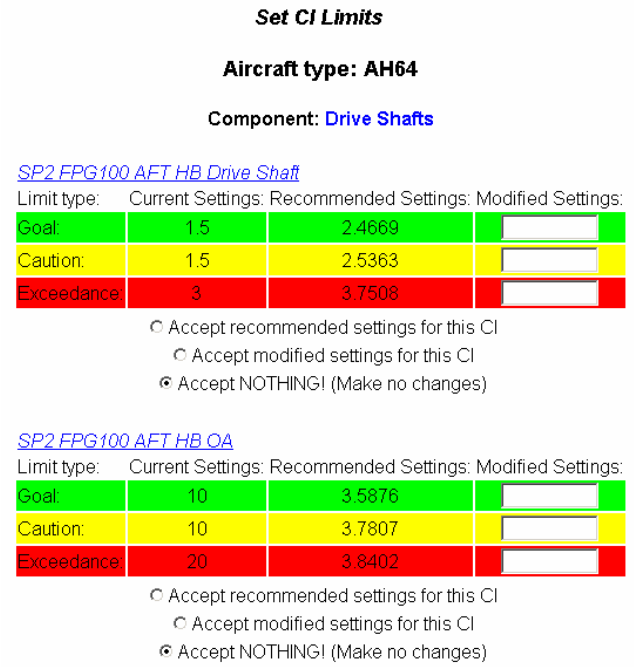


Figure 11 CI detection threshold setting

To gain more insight into the data the user may bring up additional plots for the data. Figure 12 shows some of the additional plots for the first CI for the Drive Shaft. The top shows the boxplot for the original data. The bottom plot is a histogram of that data. As described previously the boxplot gives an indication of the self consistency of the data, the spread of the data, and a poor man’s detector of outliers. All the data that was used in plots of Figure 12 had previously been labeled as “good”.

The box plots contain sets up upside down and right side up triangles. The triangles are colored green, yellow, and red and indicate the current threshold setting (the upside down triangles that appear above the center line in the plots) and the automatically generated thresholds (the right side up triangles that appear below the center line in the plots). For this particular CI as seen all of the recommended threshold settings are to the right of the current threshold settings. It is clear from this plot why false alarms would likely be occurring.

For the second CI shown in Figure 11 (SP2 FPG100 Aft HB OA) we can see that the recommended settings are much lower then the current settings. Thus with the current settings there will be no false alarms, but also no detections! Again the recommended settings would greatly improve the system’s performance.

CI Limits

Aircraft: AH64

Component: Drive Shafts

CI Name: SP2 FPG100 AFT HB Drive Shaft

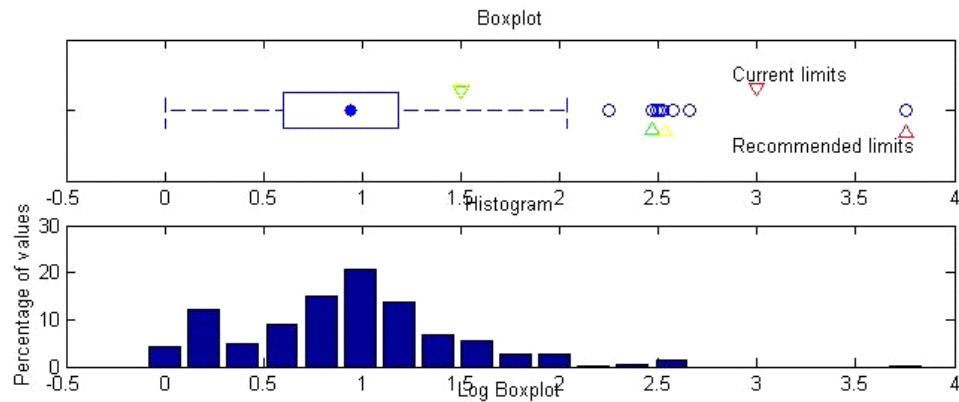


Figure 12 CI automated detection threshold

Gold Standard Data Set Development

As one can imagine, building models from examples of real data, requires that the real data be representative of the normal and fault classes that are being considered. Miss labeling of data or allowing outlier points to be included as “good” points, results in models that will not be good. Thus tools have been included in the website for viewing and including new data as it enters the system to be labeled as ‘normal’ or assigning a specific fault class. This is used in forming what we’ve called the *gold standard* data set. All new data entering the system must be reviewed by qualified personnel before being included in the gold standard database.

The Server has tools built into it to be able to bring up new data for individual components, display that data overlaid with current gold standard data. Some of those tools are described here. The data and results presented so far have been on data that has been previously reviewed and labeled as “good” or not. That same data is presented below, but before it has been labeled.

When data first enters the system, it is labeled “new.” It is not known if the data is “normal” or has some known fault condition, or is just a bad data point (i.e. equipment turned off so no vibrations generated for example). It is up to the engineer to apply a label to the data. Figure 13 shows a screen for reviewing and assigning as “good”

new data that enters into the system. The data is the AH64 Engine #2 data that has been examined previously.

In the display shown the data has been rank ordered according to its *distance* [5] from a single multi-dimensional Gaussian model fit to a transformed version of the multi-dimensional CI data. The transform converts the data from its original distribution, so that it is Gaussian in each of its CI dimensions [6]. This transformation ensures a better model fit to the data. An example of original data and transformed data is shown in Figure 15. Ensuring that the data and the form of the model fit are in agreement is often overlooked and “bad” models usually result. That transformation becomes part of the modeling information and used for all future processing (i.e. the data is first transformed, the anomaly detection of other processing applied, and the results are transformed back into the original domain). Note at the bottom of the display, the data can be sorted in a variety of other ways.

Individual points can be accepted or rejected. However a more automated approach is to use the “Rejection threshold” shown at the bottom of the display. Figure 14 shows the results of a roughly 10% rejection rate; that is that the top 10% of the points that are farthest from the nominal model are rejected. Those points are colored red in the plot; accepted points are green.

Modify Class
Aircraft: AH64
Component: Engine 2

New points

Tail Number	SP2 FPG100 #2 Eng GG	SP2 FPG100 #2 Eng PT	SP4 FPG100 #2 ENG OA	Date/Time	Status
86-08996 <input checked="" type="checkbox"/>	0.427637	0.327462	4.866620	03-Nov-2001 10:33:41	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-08996 <input checked="" type="checkbox"/>	0.411514	0.327072	4.683140	02-Nov-2001 13:15:15	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-08996 <input checked="" type="checkbox"/>	0.395290	0.337388	4.498500	03-Nov-2001 11:00:19	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-09031 <input checked="" type="checkbox"/>	0.000036	0.000399	0.000724	05-Jan-1999 00:40:21	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-09036 <input checked="" type="checkbox"/>	0.053707	0.329181	5.077570	23-May-2002 14:16:42	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
⋮					
86-08992 <input checked="" type="checkbox"/>	0.068153	0.038869	0.775602	02-May-2001 09:25:42	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-08992 <input checked="" type="checkbox"/>	0.087866	0.125816	0.014492	01-Apr-2002 11:57:23	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-08992 <input checked="" type="checkbox"/>	0.081643	0.159876	0.011818	01-Apr-2002 12:26:20	<input checked="" type="radio"/> Accept <input type="radio"/> Reject
86-08992 <input checked="" type="checkbox"/>	0.072278	0.062366	0.010277	11-Apr-2002 09:12:42	<input checked="" type="radio"/> Accept <input type="radio"/> Reject

☒ Accept changes above
 ☐ Accept all
 ☐ Reject all
 ☐ Rejection threshold:

Sort by: (Saves changes to status column)

☐ Tail Number
☐ SP2 FPG100 #2 Eng GG
☐ SP2 FPG100 #2 Eng PT
☐ SP4 FPG100 #2 ENG OA
☐ Date/Time
☒ Value

Figure 13 Golden data set specification

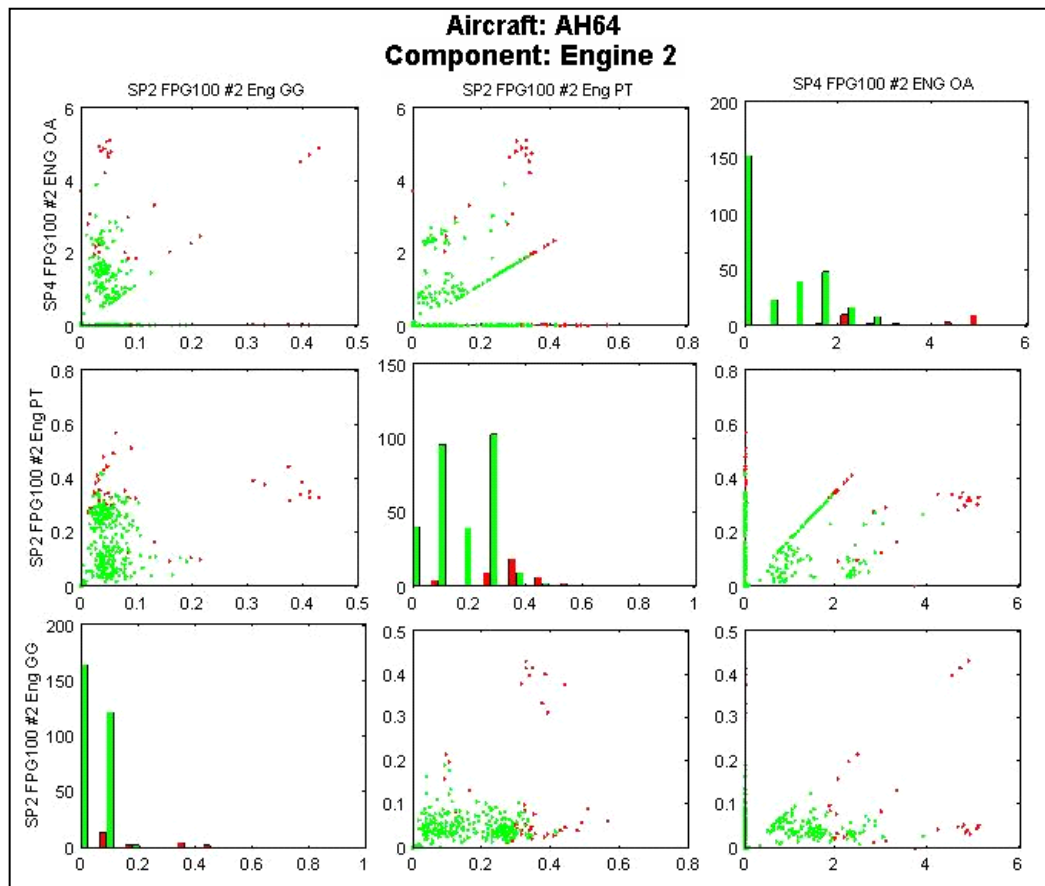


Figure 14 Labeled data




Figure 15 Gaussian transformation example

The histograms along the diagonals contain red and green bars that correspond to the red and green points. At this point the user can blow up individual displays, highlight single points to determine which tail number / flight that data was collected from, and then pull up the data as originally collected for even more detailed review.

Once the data has been validated, the user pushes a button to accept the labeling, and those new samples will be included into the gold standard database. As more samples are collected, the automated portion of the processing will become more reliable.

Data in the tails of the normal vs. fault distributions will start to be filled in!

5. SUMMARY AND RECOMMENDATIONS

IAC, in collaboration with the US Army and Air Force is developing a system for monitoring of Army helicopter vibration, engine performance, and aircraft structures. The processing being developed is part of the Army's Vibration Management Enhancement Program (VMEP).

VMEP is composed of three primary components; (1) an on-board system, (2) a ground-based station, and (3) a web server. Here we've concentrated on the VMEP Web Server component. The Server is used to collect and process data from all Army helicopters that have the VMEP on-board system.

Additions to the system not discussed here include engine performance data collection / monitoring and flight regime recognition in support of structural life monitoring and automated triggering of vibration data collections. Almost all of the tools developed for processing data are based on the Matlab and Simulink processing environment. IAC has developed an extensive set of Matlab based tools that include data import, vibration data analysis, performance and prediction models, statistical models, regime recognition, and reasoning tools.

Collection of large amounts of data is required to capture both normal and know fault events. All the action in

prognosis occurs in the tails of distributions of "normal" vs. "fault" measurements – precisely the areas where not much data has been collected. This data is required for empirical model development but also for validation of detailed physics based models. The tails are a scary place. A slight change in a threshold can mean a drastic change in times.

The VMEP system is an ideal data collection and processing testbed for the continuing collection and analysis of large amounts of real data in support of automated diagnostic and prognostic system development. The VMEP Server included tools to automatically screen "normal" and known fault data so as to detect off-nominal events that are of interest. It includes tools to systematically continuously update existing model parameters to improve overall detection and classification performance.

6. REFERENCES

- [1] T. Brotherton, P. Grabill, D. Wroblewski, R. Friend, B. Sotomayer, and J. Berry, "A Testbed for Data Fusion for Engine Diagnostics and Prognostics", *Proceedings of the 2002 IEEE Aerospace Conference*, Big Sky MT, March 2002
- [2] P. Grabill, J. Berry, L. Grant, and J. Porter, "Automated Helicopter Vibration Diagnostics for the US Army and National Guard", *57th Annual Forum of the American Helicopter Society*, Washington, DC, May 2001.
- [3] P. Grabill, T. Brotherton, J. Berry, and L. Grant, "The US Army and National Guard Vibration Management Enhancement Program (VMEP): Data Analysis and Statistical Results", *American Helicopter Society 58th Annual Forum*, Montreal, Canada, June 10-12, 2002.
- [4] T. Brotherton and R. Mackey, "Anomaly Detector Fusion Processing for Advanced Military Aircraft", *The 2001 IEEE Aerospace Conference*, Big Sky, MT, March 2001.
- [5] T. Brotherton and T. Johnson, "Anomaly Detection for Advanced Military Aircraft using Neural Networks", *The 2001 IEEE Aerospace Conference*, Big Sky, MT, March 2001.
- [6] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, 1965.



Tom Brotherton received his B.S. degree from Cornell University (1974), an M.A.Sc. from the University of Toronto (1976) and the Ph.D. from the University of Hawaii (1982) all in electrical engineering. He was an assistant professor in the Information and Computer Sciences Department at the University of Hawaii in 1983 and with the Orincon Corp from 1983-1999. He is currently a VP of the Intelligent Automation Corp. (IAC) in San Diego, CA. IAC is involved with the development of aircraft and related equipment monitoring software and hardware. Dr. Brotherton is also on the editorial board for the IEEE Press series on Biomedical Engineering. His interests are in the development of adaptive signal and data fusion techniques for machine condition, fault monitoring, and prognostics.



Paul Grabill received his B.S. degree in mechanical engineering in 1986 at the University of Cincinnati. He worked for 10 years at General Electric Aircraft Engines as a dynamics engineer where he was involved in gas turbine vibration testing, engine balancing, and system dynamic analysis. He received his M.S. degree in mechanical engineering in 1992 at the University of Cincinnati where he was sponsored by the NASA Health Monitoring Center and UC's Structural Dynamics Research Laboratory. He is currently Vice President of Engineering at Intelligent Automation Corporation in San Diego where he is currently developing advanced diagnostic and prognostic systems for helicopters, gas turbines, and aircraft.



Wing Commander Richard Friend is a Royal Air Force engineering officer. He was serving on an exchange program in the Propulsion Directorate at the United States Air Force Research Laboratory, Wright Patterson Air Force Base, Ohio. There he was the research group leader for USAF Engine Health Management and Control, the Secretariat for the Steering Committee for the High Cycle Fatigue Science and Technology Initiative, and was the USAF lead for the Joint Strike Fighter engine Prognostic Health Management systems. An aeromechanical engineer, he is specialized in gas turbine engines and has accrued over 24-years experience in the maintenance and management of fighter aircraft, engines and engine test facilities.

No
picture
available

Bill Sotomayer was with the Propulsion Directorate at the United States Air Force Research Laboratory, Wright Patterson Air Force Base, Ohio. There he was actively involved in Engine Health Management (EHM) programs in support of Air Force C17, the Joint Strike Fighter, and the USAF's research thrust, the Versatile Affordable Advanced Turbine Engine; VAATE.. Previously, he was involved in friction modeling for blade/disk interfaces on high performance jet engines.



John Berry received his B.S., M.S. and Ph.D. degrees from the Georgia Institute of Technology (1974, 78, 90), all in aerospace engineering. He served for 8 years as a signal officer in the US Army. From 1982 until 1998, Dr. Berry served the US Army as a civilian research engineer in experimental and analytical rotorcraft aerodynamics at the NASA Langley Research Center. Since 1998 he has been an aeromechanics branch chief for the Aviation and Missile Research, Development, and Engineering Center. Dr. Berry is a past associate-editor for the Journal of the American Helicopter Society and contributes actively to that professional society. His interests are in the development of innovative technologies that contribute to improved rotorcraft aeromechanical systems.